arXiv:1604.04741v1 [stat.ML] 16 Apr 2016

Smoothed Hierarchical Dirichlet Process: A Non-Parametric Approach to Constraint Measures

Cheng Luo, Yang Xiang and Richard Yi Da Xu

April 19, 2016

Abstract

Time-varying mixture densities occur in many scenarios, for example, the distributions of keywords that appear in publications may evolve from year to year, video frame features associated with multiple targets may evolve in a sequence. Any models that realistically cater to this phenomenon must exhibit two important properties: the underlying mixture densities must have an unknown number of mixtures; and there must be some "smoothness" constraints in place for the adjacent mixture densities. The traditional Hierarchical Dirichlet Process (HDP) may be suited to the first property, but certainly not the second. This is due to how each random measure in the lower hierarchies is sampled independent of each other and hence does not facilitate any temporal correlations. To overcome such shortcomings, we proposed a new Smoothed Hierarchical Dirichlet Process (sHDP). The key novelty of this model is that we place a temporal constraint amongst the nearby discrete measures $\{G_j\}$ in the form of symmetric Kullback-Leibler (KL) Divergence with a fixed bound B. Although the constraint we place only involves a single scalar value, it nonetheless allows for flexibility in the corresponding successive measures. Remarkably, it also led us to infer the model within the stick-breaking process where the traditional Beta distribution used in stick-breaking is now replaced by a new constraint calculated from B. We present the inference algorithm and elaborate on its solutions. Our experiment using NIPS keywords has shown the desirable effect of the model.

Key-words: Bayesian non-parametric, smoothed HDP, Bayesian inference, particle filtering, truncated Beta distribution.

1 Introduction

The Hierarchical Dirichlet Process (HDP) is an extension to the traditional Dirichlet process (DP), which essentially is comprised of a set of measures responsible for generating observations in each of its respective groups. The model allows the lower hierarchies of different groups to share "atoms" of their parents. They allow practitioners to apply the model to scenarios where independent mixture densities (of each group) share certain mixture components amongst its siblings which are inherited from their parents.

HDP is proposed by [17] and further studied by [18][19][14][8] and many other literatures in various disciplines have used it as an infinite version of Latent Dirichlet Allocation (LDA) [3]. In order to perform inference, [17] proposed three Monte-Carlo Markov Chain (MCMC) methods, two of which adopted collapsed Gibbs with G_0 and G_j integrated out while the third one instantiated G_0 and integrated G_j. The other works focus on the variational inference of HDP.

HDP has been successfully extended to model sequential data, in the form of so called HDP-HMM [17][19][8] and infinite HMM [2]. In these models, it is assumed that there exists a set of "time invariant measures"; A series of latent states are then drawn from these time invariant measures. The index for which a measure is to be used at time t, is determined by the state of the previous state, i.e., $p(z_t|z_{t-1}, G_1, \dots G_k) = G_{z_{t-1}}$. In order to cater for a "smooth" transition of states, [7] proposed the so-called 'sticky' HDP-HMM method. This approach adds a weight for the self-transition bias and places a separate prior on this prior. While such methods can be simpler in terms of inference, and may make more sense in situations where the states distributions may repeat at some point in time, they are not suitable for scenarios where the underlying distribution does not repeat in cycles, i.e., there is a need for a unique distribution to be assigned

at each time t. One example of this is the distribution of topics for a scientific journal, it is highly unlikely that its topic distributions will "come" back at some years in the future.

In order to generate a set of non-repeating and time-varying measures associated with observations at time t, [12] proposed a Bayesian non-parametric approach in the form of Dependent Dirichlet processes. Based on an initial Dirichlet process, successive measures can be formed through three operations: superposition, subsampling and point transition. However, each of its operations can be very restrictive. For example, in sub-sampling, the "next" measure must have a fewer number of "sticks" than the first. The full flexibilities can only be achieved using combinations of all three operations, which makes the method very complex.

In this paper, we propose an alternative method to construct a time-varying dependent Dirichlet processes. Our model uses a very simple constraint in which the "smoothness" of the mixing measures in the second layer of HDP is achieved by simply placing a symmetric KL divergence constraint between them which is bounded by some fixed value B. The key motivation is achieved through the following observation: Using our KL constraint, we can achieve our intended outcome by substituting the Beta distribution with a truncated Beta while still using the stick-breaking paradigm. The new truncation is calculated from any B valued placed, and we observed that there are only a finite number of possible solution spaces for the truncations. Subsequently, we developed a sampling method using a Gibbs Sampling framework, where one of the Gibbs steps is achieved using particle filtering. We named our method Smoothed HDP (sHDP). Since we only applied a simple scalar value for its constraint, we argue it is the non-parametric approach to place constraints amongst Dirichlet Processes.

This rest of the paper is organised as follows: The second section describes the background of our model, and for completeness we include the related inference method. The third section is the elaborated description of our model. In section 4 we give the details of the inference method. In section 5, we test the model using synthetic data sets and the PAMI article key words. Section 6 provides a conclusion and some further discussions.

2 Background

To make this paper self-contained, we describe in this paper, some of the related models and their associated inference methods. We adopt the stick breaking paradigm for the representation of a Dirichlet process for both G_0 and G_j. In terms of inference, given both G_0 and G_j, particle filtering is used instead of Gibbs sampling to infer G_{j+1} since the conditional distribution, $\mathbb{P}(G_{j+1}|G_0, G_j)$ cannot be obtained analytically. We use the slice sampling proposed by [4] for the sampling of the truncated Beta distribution.

2.1 The hierarchical Dirichlet process

Let H be a diffusion measure (a measure having no atoms a.s.) [11] and γ be a scalar, and a discrete measure G_0 is sampled from $\mathrm{DP}(\gamma, H)$ [16]. Setting G_0 as the base measure and $\alpha > 0$ be the concentration parameter, we sample $G_j \sim \mathrm{DP}(\alpha, G_0)$ for $j = 1, ..., M$. Obviously, G_j should be discrete and some of the activated atoms for different G_j are shared. We the adopt stick-breaking paradigm for the representation of a Dirichlet process, and denote

$$G_0 = \sum_{i=1}^{\infty} \pi_i \delta_{\phi_i}, \tag{1}$$

where δ is the Dirac function,

$$\tilde{\pi}_i \overset{i.i.d.}{\sim} \mathrm{Beta}(1, \gamma), \quad \pi_i = \tilde{\pi}_i \prod_{l=1}^{i-1} \tilde{\pi}_l,$$

$$\phi_i \overset{i.i.d.}{\sim} H,$$

and $\tilde{\pi}_i$ and ϕ_i are mutually independent. As proposed by [17], G_j can be formulated as

$$G_j = \sum_{i=1}^{\infty} \beta_i \delta_{\phi_i}, \quad \text{for } j = 1, ..., M, \tag{2}$$

where

$$\tilde{\beta}_i \sim \text{Beta}\left(\alpha\pi_i, \alpha\left(1 - \sum_{l=1}^{i}\pi_l\right)\right), \quad \beta_i = \tilde{\beta}_i \prod_{l=1}^{i-1}(1-\tilde{\beta}_l). \tag{3}$$

This formulation has an advantage in that the atoms in (2) are distinct.

2.2 Posterior sampling of the GEM distribution

The Chinese restaurant process is famous in the posterior sampling of the Dirichlet process. However, since our model requires to sample G_0 and G_j explicitly, we use the method proposed by [9]. Denoting by m_i the number of observations equal to i, the posterior of G_0 is drawn as

$$\begin{aligned} \tilde{\pi}_i &\sim \text{Beta}\left(1 + m_i, \gamma + \sum_{l=i+1}^{\infty} m_l\right), \\ \pi_i &= \tilde{\pi}_i \prod_{l=1}^{i-1}(1-\tilde{\pi}_l). \end{aligned} \tag{4}$$

2.3 Slice sampling of a truncated Beta distribution

For the sampling of the truncated Beta distribution, we use slice sampling proposed by [4]. Let $0 \le a < b \le 1$ and $\alpha, \beta > 0$ be scalars and the density of X is

$$f(x|a,b,\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}\delta(a < x < b),$$

where $\delta(\cdot)$ is the Dirac function.To sample from this truncated Beta distribution, we add auxiliary variable U and construct a joint density of X, U as

$$p(x,u) \propto x^{\alpha-1}\delta(0 < u < (1-x)^{\beta-1}, a < x < b).$$

The method of drawing of x and u is:

1. Sampling u conditional on x
$$u \sim \text{Uniform}(0, (1-x)^{\beta-1}).$$
2. Sampling x conditional on u using inverse transform
   (a) if $\beta > 1$
   $$p(x) \propto x^{\alpha-1}\delta(a < x < \min\{1 - u^{1/(\beta-1)}, b\}),$$
   (b) if $\beta < 1$
   $$p(x) \propto x^{\alpha-1}\delta(\max\{1 - u^{1/(\beta-1)}, a\} < x < b).$$

2.4 Particle filtering

Particle filtering is used to sample the posterior distribution of sequential random variables. The advantage of particle filtering is that sampling is possible when the posterior density cannot be stated analytically. Let $\{z_t : t = 1, ..., T\}$ be latent random variables and $\{x_t : t = 1, ..., T\}$ be observed random variables. The joint distribution of them is

$$p(x_{1:t}, z_{1:t}) \propto p(z_1) \prod_{t=1}^{T-1} p(z_{t+1}|z_t) \prod_{t=1}^{T} p(x_t|z_t).$$

Let N be a positive integer large enough, and the posterior sampling method is:

1. For $i = 1, ..., N$, propose $z_1^{(i)}$ from prior $p(z_1)$, and compute weights $w^{(i)} = p(x_1|z_1)$.

2. Resample $\boldsymbol{z}_1$ with weights $\boldsymbol{w}$, where $\boldsymbol{z}_1$ denotes the vector of $(z_1^{(i)})_{i=1,...,N}$ and $\boldsymbol{w}$ denotes $(w^{(i)})_{i=1,...,N}$.

3. For $t = 2, ..., T$,

   (a) propose $z_t^{(i)}$ from $p(z_{t+1}|z_t)$, and compute weights $w^{(i)} = p(x_t|z_t)$.

   (b) resample $\boldsymbol{z}_t$ with weights $\boldsymbol{w}$.

The empirical posterior distribution of $(z_t)_{t=1,...,T}$ is

$$p(z_{1:T}) = \frac{1}{N}\sum_{i=1}^{N} \delta_{z_{1:T}^{(i)}}. \tag{5}$$

For a fully description of particle filtering please refer to [5].

3 The model of the smoothed HDP

In this section, we will give a full description of our model. In the traditional HDP, the discrete measures in the lower level, i.e., G_j, are independent given the concentration parameter α and the base measure G_0. Separately from this independent assumption, we force the successive mixing measures G_j and G_{j+1} to be alike in terms of their symmetric KL divergence. Since both G_j and G_{j+1} are discrete measures with infinite many atoms, the computation of symmetric KL divergence between them is intractable. As a substitution, we propose an aggregated form of symmetric KL divergence that is computable. It can be proved that our aggregated symmetric KL divergence has the same expectation as the original form although they are not equal all the time.

3.1 The aggregated symmetric KL divergence

Let $\gamma, \alpha > 0$ be scalars and H be a diffusion measure, we sample a discrete base measure G_0 distributed as $\mathrm{DP}(\gamma, H)$. Then the mixing measures (G_j) in the lower level are sampled from

$$G_j \sim \mathrm{DP}(\alpha, G_0), \quad \text{for } j = 1, ..., M,$$

with a constraint that the symmetric KL divergence $\mathrm{KL}(G_{j+1}||G_j) < B$ for some fixed positive scalar B. Suppose G_0 and G_j have the following expression

$$G_0 = \sum_{i=1}^{\infty} \pi_i \delta_{\phi_i}, \quad G_j = \sum_{i=1}^{\infty} \beta_{j,i} \delta_{\phi_i}.$$

The symmetric KL divergence between G_j and G_{j+1} is defined to be

$$\mathrm{KL}(G_j||G_{j+1}) = \sum_{i=1}^{\infty} \beta_{j,i} \log \frac{\beta_{j,i}}{\beta_{j+1,i}} + \sum_{i=1}^{\infty} \beta_{j+1,i} \log \frac{\beta_{j+1,i}}{\beta_{j,i}}. \tag{6}$$

Our problem is summarised as

Problem 1. *Given G_0 and G_j,*

$$\begin{aligned} &\textit{sample} && G_{j+1} \sim \mathrm{DP}(\alpha, G_0), \\ &\textit{subject to} && \mathrm{KL}(G_j||G_{j+1}) \leq B, \end{aligned}$$

where $B > 0$ is a scalar.

However, since (6) is a sum of infinite terms, the direct computation of (6) is intractable. Alternatively, for every atom l, we define an aggregated symmetric KL divergence as

$$\begin{aligned}
&\text{aggKL}(l; G_j||G_{j+1}) \\
&= \left(\sum_{i=1}^{l-1} \beta_{j,i}\right) \log \frac{\sum_{i=1}^{l-1} \beta_{j,i}}{\sum_{i=1}^{l-1} \beta_{j+1,i}} \\
&\quad + \beta_{j,l} \log \frac{\beta_{j,l}}{\beta_{j+1,l}} + \left(\sum_{i=l+1}^{\infty} \beta_{j,i}\right) \log \frac{\sum_{i=l+1}^{\infty} \beta_{j,i}}{\sum_{i=l+1}^{\infty} \beta_{j+1,i}} \\
&\quad + \left(\sum_{i=1}^{l-1} \beta_{j+1,i}\right) \log \frac{\sum_{i=1}^{l-1} \beta_{j+1,i}}{\sum_{i=1}^{l-1} \beta_{j,i}} + \beta_{j+1,l} \log \frac{\beta_{j+1,l}}{\beta_{j,l}} \\
&\quad + \left(\sum_{i=l+1}^{\infty} \beta_{j+1,i}\right) \log \frac{\sum_{i=l+1}^{\infty} \beta_{j+1,i}}{\sum_{i=l+1}^{\infty} \beta_{j,i}}.
\end{aligned} \tag{7}$$

Hence Problem 1 changes to

Problem 2. *Given G_0 and G_j, for $l = 1, 2, 3, ...$*

$$\begin{aligned}
&\textit{sample} && G_{j+1} \sim \text{DP}(\alpha, G_0), \\
&\textit{subject to} && \text{aggKL}(l; G_j||G_{j+1}) \le B^*,
\end{aligned} \tag{8}$$

where $B^* > 0$ is a scalar.

The simplest way to set the value of B^* is letting $B^* = B$. But this does not make sense unless $\text{KL}(G_j||G_{j+1}) = \text{aggKL}(G_j||G_{j+1})$. However, direct algebraic calculation shows that the these two symmetric KL divergences are not agree at all times. Fortunately, the next theorem shows that the two terms, $\text{KL}(G_j||G_{j+1})$ and $\text{aggKL}(G_j||G_{j+1})$, take the same value on average.

Theorem 1. *The expectation of $\text{KL}(G_j||G_{j+1})$ and $\text{aggKL}(l; G_j||G_{j+1})$ agrees with respect to $\boldsymbol{\beta}_j, \boldsymbol{\beta}_{j+1}$, where $\boldsymbol{\beta}_j = (\beta_{j,i})_{i=1,2,...}$ and $\boldsymbol{\beta}_{j+1} = (\beta_{j+1,i})_{i=1,2,...}$.*

Proof. Subtract (7) from (6) gives

$$\begin{aligned}
&\text{KL}(G_j||G_{j+1}) - \text{aggKL}(l; G_j||G_{j+1}) \\
&= \sum_{i=1}^{l-1} \left(\beta_{j,i} \log \frac{\beta_{j,i}}{\sum_{i=1}^{l-1} \beta_{j,i}} - \beta_{j+1,i} \log \frac{\beta_{j+1,i}}{\sum_{i=1}^{l-1} \beta_{j+1,i}}\right) \\
&\quad + \sum_{i=l+1}^{\infty} \left(\beta_{j,i} \log \frac{\beta_{j,i}}{\sum_{i=l+1}^{\infty} \beta_{j,i}} - \beta_{j+1,i} \log \frac{\beta_{j+1,i}}{\sum_{i=l+1}^{\infty} \beta_{j+1,i}}\right).
\end{aligned} \tag{9}$$

Since $\beta_{j,i}$ and $\beta_{j+1,i}$ has the same distribution for all $i = 1, 2, 3, ...$, (9) must has expectation 0 provides $\log \sum_{i=l+1}^{\infty} \beta_{j,i}$ and $\sum_{i=l+1}^{\infty} \beta_{j+1,i}$ are finite. But this is obvious since $\sum_{i=l+1}^{\infty} \beta_{j+1,i} = 1 - \sum_{i=1}^{l} \beta_{j,i}$ and $\sum_{i=l+1}^{\infty} \beta_{j+1,i} = 1 - \sum_{i=1}^{l} \beta_{j+1,i}$. □

In light of Theorem 1, we are safe to set $B = B^*$. The details of sampling G_{j+1} is given in the next subsection.

3.2 How to sample G_{j+1} conditional on G_j

Suppose G_0, G_j and $\beta_{j+1,1:l-1}$ is known, and we want to sample $\beta_{j+1,l}$. Observe (7) and with a few algebraic calculation, the aggregated symmetric KL divergence constraint can be restated as

$$\begin{aligned}
&- \left(1 - \sum_{i=1}^{l} \beta_{j,i}\right) \log \left(1 - \sum_{i=1}^{l-1} \beta_{j+1,i} - \beta_{j+1,l}\right) \\
&+ \left(1 - \sum_{i=1}^{l-1} \beta_{j+1,l} - \beta_{j+1,i}\right) \log \frac{1 - \sum_{i=1}^{l-1} \beta_{j+1,i} - \beta_{j+1,l}}{\sum_{i=l+1}^{\infty} \beta_{j,i}} \\
&\beta_{j+1,l} \log \frac{\beta_{j+1,l}}{\beta_{j,l}} - \beta_{j,l} \log \beta_{j+1,l} \le B - C,
\end{aligned} \tag{10}$$

where C is a constant with respect to $\beta_{j+1,l}$, which is

$$C = \left(\sum_{i=1}^{l-1} \beta_{j,i}\right) \log \frac{\sum_{i=1}^{l-1} \beta_{j,i}}{\sum_{i=1}^{l-1} \beta_{j+1,i}} + \beta_{j,l} \log \beta_{j,l} + \left(\sum_{i=1}^{l-1} \beta_{j+1,i}\right) \log \frac{\sum_{i=1}^{l-1} \beta_{j+1,i}}{\sum_{i=1}^{l-1} \beta_{j,i}}.$$

Note the function

$$f(x) = x \log \frac{x}{\beta_{j,l}} - \beta_{j,l} \log x - \left(1 - \sum_{i=1}^{l} \beta_{j,i}\right) \log \left(1 - \sum_{i=1}^{l-1} \beta_{j+1,i} - x\right) + \left(1 - \sum_{i=1}^{l-1} \beta_{j+1,l} - x\right) \log \frac{1 - \sum_{i=1}^{l-1} \beta_{j+1,i} - x}{\sum_{i=l+1}^{\infty} \beta_{j,i}}$$

is convex in the interval $\left(0, 1 - \sum_{i=1}^{l-1} \beta_{j+1,i}\right)$, the equality

$$f(x) = B - C \tag{11}$$

has at most two roots. Showing in Figure 1, inequality (10) has solution of form

- $\beta_{j+1,l} \in (r_1, r_2)$ if (11) has two solutions r_1 and r_2,
- $\beta_{j+1,l} \in (0, r_1)$ or $(r_1, 1)$ if (11) has only one solution r_1.
- If (11) has no roots, either the entire interval $\left(0, 1 - \sum_{i=1}^{l-1} \beta_{j+1,i}\right)$ is the solution or B is too tight that there is no discrete measure satisfies condition (8).

According to the relationship of $\beta_{j+1,l}$ and $\tilde{\beta}_{j+1,l}$ shown in (2), the truncating boundary for $\tilde{\beta}_{j+1,l}$ are $r_1 / \prod_{i=1}^{l-1}(1 - \tilde{\beta}_{j+1,i})$ and $r_2 / \prod_{i=1}^{l-1}(1 - \tilde{\beta}_{j+1,i})$, where r_1 is the lower bound and r_2 is the upper bound and they can be 0 or 1. Combine this constraint and the sampling method of β_j in (2)(3) gives the solution of Problem 2 in Algorithm 1.

Algorithm 1 Sampling G_{j+1} with aggregated symmetric KL divergence constraint.

Input: Base measure G_0, predecessor G_j, concentration parameter α and the symmetric KL divergence bound B.

Output: A discrete measure G_{j+1} satisfying condition (8).

1: **for** $l = 1, 2, 3, ...$ **do**
2: Derive the roots of (11).
3: Compute truncating boundary for $\tilde{\beta}_{j+1,l}$.
4: Sample $\tilde{\beta}_{j+1,l}$ from a truncated Beta distribution.
5: **end for**
6: **for** l = 1,2,3,... **do**
7: Compute $\beta_{j+1,l}$.
8: **end for**

4 Inference

Note in Algorithm 1 we sample G_{j+1} directly without the explicit likelihood, hence we cannot compute the posterior of G_{j+1} analytically. Consequently, we have to use particle filtering to sample the posterior distribution of G_j for $j = 1, ..., M$.

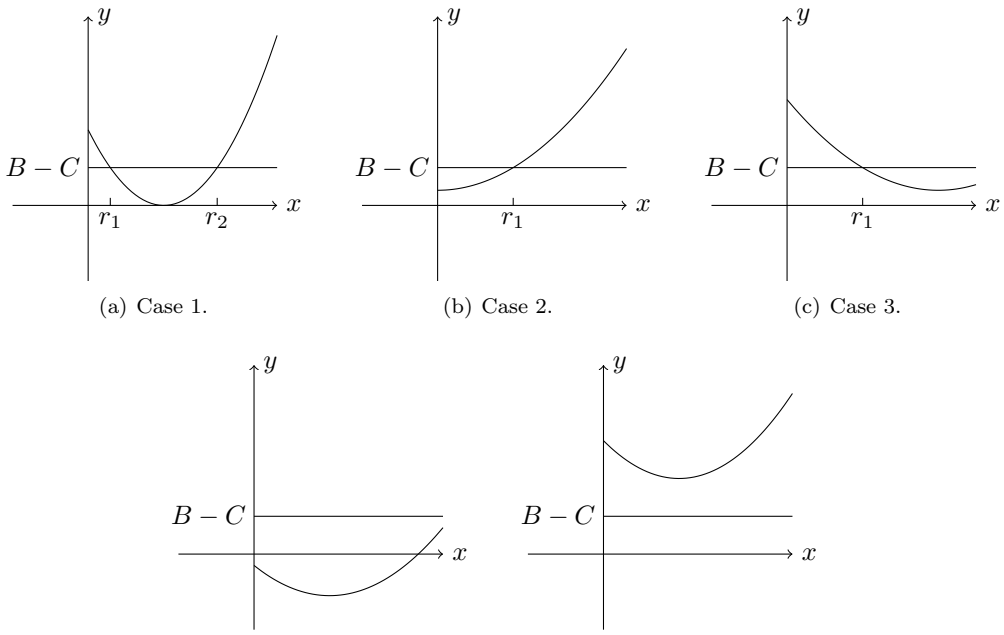

(a) Case 1. (b) Case 2. (c) Case 3.

(d) Case 4. (e) Case 5.

Figure 1: The cases of solution (10). The feasible interval for case (a)-(e) is (r_1, r_2), $(0, r_1)$, $(r_1, 1)$, $(0, 1)$, $\emptyset$ respectively.

4.1 Initialization

According to the definition of Dirichlet process [6], given concentration parameter α and base measure G_0, the expectation of distributions sampled from $\mathrm{DP}(\alpha, G_0)$ is G_0. Let $d_{j,h}$ be an observation in phase j, then the distribution of $d_{j,h}$ is modeled as

$$d_{j,h} \sim G_j, \quad G_j \sim \mathrm{DP}(\alpha, G_0). \tag{12}$$

Followed by the deduction above and integrate out G_j, the distribution of $d_{j,h}$ is

$$d_{j,h} \sim \int G_j(d_{j,h}) dG_j = G_0. \tag{13}$$

Hence the posterior distribution of G_0 can be sampled with help of (4). Suppose the prior of G_0 is a Dirichlet process $\mathrm{DP}(\gamma, H)$ and G_0 has the form of (1). Let c_i be a set of the observations taking on cluster i and $m_i = |c_i|$, then the posterior of $\tilde{\pi}_i$ is

$$p(\tilde{\pi}_i | c_i, c_{i+1}, ..., \gamma) \sim \mathrm{Beta}\left(1 + m_i, \gamma + \sum_{l=i+1}^{\infty} m_l\right),$$

and the posterior of π_i is

$$\pi_i = \tilde{\pi}_i \prod_{l=1}^{i-1} (1 - \tilde{\pi}_l).$$

The posterior of ϕ_i is

$$p(\phi | c_i, d_{j,\cdot}) \propto H(\phi_i) \prod_{h \in c_i} F(d_{j,h} | \phi_i),$$

where $d_{j,\cdot}$ denote all the observations in phase j and $F(\cdot)$ is the likelihood function.

4.2 Update of G_j

Given G_0, G_j and α, we use particle filtering to sample the posterior of G_j. With the help of Algorithm 1, we give a proposal G_{j+1}. The weight of this proposal is a product of probabilities

$$w = \prod_i \beta_i^{m_{j,i}}, \tag{14}$$

where $m_{j,i}$ is the number of observations in phase j that take cluster i. Using particle filtering, our sampling algorithm of $G_{1:M}$ is given in Algorithm 2.

Algorithm 2 Particle filtering for $G_{1:M}$

Input: The base measure G_0, concentration parameter α, symmetric KL divergence bound B, observations d and a large integer N.

Output: Discrete measures $G_{1:M}$ satisfying condition (8) for each pair of successive G_j and G_{j+1}.

1: For $n = 1, ..., N$, sample G_1 by (2)(3).
2: **for** $j = 2, ..., M$ **do**
3: For $n = 1, ..., N$, sample $G_j^{(n)}$ with Algorithm 1.
4: Compute importance weights $w_j^{(n)}$ by (14).
5: Normalize importance wights.
6: Resample $G_j^{(n)}$ with importance weights.
7: **end for**

4.3 Update of z

The latent variable $z_{j,h}$ is used to indicate which cluster the observation $d_{j,h}$ is from. By the definition of G_j, the probability of $z_{j,h} = i$ is $\beta_{j,i}$. Given $F(\cdot)$ and ϕ_i, the likelihood is $F(d_{j,h}|\phi_i)$. Combining these together gives the posterior probability of $z_{j,h}$, which is proportional to

$$p(z_{j,h} = i|\phi_i, \beta_{j,i}) \propto \beta_{j,i} F(d_{j,i}|\phi_i),$$

where $\beta_{j,i}$ is the weight of the atom ϕ_i in G_j.

4.4 Update of the hyper-parameters γ, α

For hyper-parameters γ and α, we simply put Gamma distribution as priors on them, namely, Gamma(a_γ, b_γ) and Gamma(a_α, b_α). The posterior of γ and α are

$$p(\gamma|a_\gamma, b_\gamma, k, m) \propto \gamma^{a_\gamma - 2} e^{-b_\gamma \gamma} (\gamma + m) \int_0^1 u^\gamma (1-u)^{m-1} du,$$
$$p(\alpha|a_\alpha, b_\alpha, k, m) \propto \alpha^{a_\alpha - 2} e^{-b_\alpha \alpha} (\alpha + m) \int_0^1 u^\alpha (1-u)^{m-1} du.$$

respectively, where m is the number of observations and k is the number of distinct atoms. By adding auxiliary variable u, it renders a mixture of Gammas for γ and α. For details of the sampling method please refer to [20].

5 Experiments

We validate our smoothed HDP model using both synthetic and real dataset. MATLAB code of our experiments are available in the following website: https://github.com/llcc402/MATLAB-codes.

We assume the number of activated atoms is at most $K = 100$, hence the definition of the symmetric KL divergence (6) becomes:

$$\mathrm{KL}(G_j||G_{j+1}) = \sum_{i=1}^K \beta_{j,i} \log \frac{\beta_{j,i}}{\beta_{j+1,i}} + \sum_{i=1}^K \beta_{j+1,i} \log \frac{\beta_{j+1,i}}{\beta_{j,i}}. \tag{15}$$

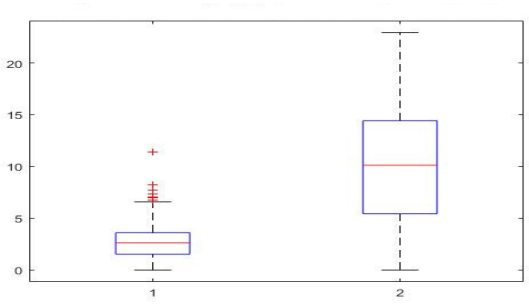

(a) A comparison of the symmetric KL divergence for sHDP and HDP. Left is the smoothed sampling result and right is the traditional sampling result.

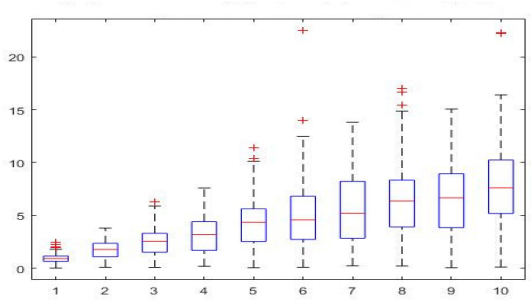

(b) The KL divergences sHDP when the bound B changes. From left to right the bound B is set from 1 to 10.

Figure 2: Comparing the symmetric KL divergences.

In case some weights of G_j become close to zero, causing the KL to become arbitrarily large, we set a minimum lower bound for every component weight to be $\epsilon = 1\text{e} - 5$.

5.1 Simulations on synthetic data set

Since the computation of the symmetric KL divergence only requires the knowledge of weights of the discrete measures, we discard the positions and sample GEM(γ) instead. In our experiments, we set $\gamma = 5$. Conditioning on G_0, we sample $G_1 \sim \text{DP}(\alpha, G_0)$, we then sample $G_2' \sim \mathbb{P}(G_2|\alpha, G_0, G_1)$ using our sHDP. We compare our model with HDP in order to show its smoothness effect. In HDP, while having G_0 and G_1 sampled in the same way as sHDP, G_2 is instead to be generated independently, i.e., $G_2 \sim \text{DP}(\alpha, G_0)$. In all the experiments we set $\alpha = 1$.

In sHDP, we set the KL bound $B = 3$, and compare the symmetric KL divergence between G_1, G_2 and G_1, G_2'. We repeat this procedure for 1000 times and compare the symmetric KL divergence by (15), and the result is shown in Figure 2(a). From this figure, we can see that the symmetric KL divergence of successive measures are much smaller than the traditional model. The mean of the values of sHDP is about 3, just as Theorem 1 illustrated, while the mean of traditional HDP is about 10.

In the second experiment, we show how the mean of KL varies with the bound B we choose. By setting B from 1 to 10, we sample G_1, G_2' 100 times and compute the symmetric KL divergence between them. This is shown in Figure 2(b). We show that the average symmetric KL divergence of the two distributions have increased when the bound value increases, which leads to an expected outcome.

The last simulation is about using the time series data with the following setting: We first set $\gamma = 5$, $\alpha = 1$ and $B = 1$, then we generate $G_1, ..., G_{20}$ using sHDP. After that, we add Gamma noise to the distributions $G_1, ..., G_{20}$. The noise is simulated as follows. First generate $K \times 20$ independent random variables distributed as Gamma$(0.03, 1)$, then add them to $G_1, ..., G_{20}$. Lastly, we normalize the distributions. With these noisy

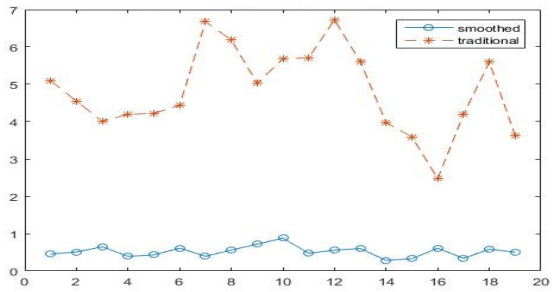

(a) The symmetric KL divergence for successive distributions. The blue solid line is the values of the KL divergence for sHDP, and the red dotted line is the values of the KL divergence for HDP.

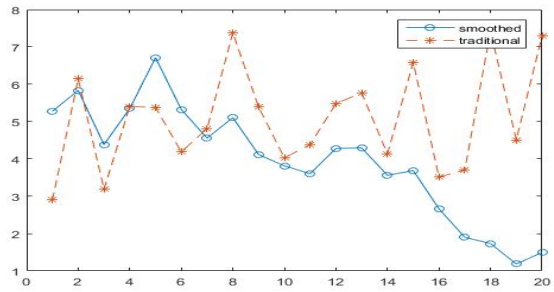

(b) The distance between theoretic value $G_{1:20}$ with the posterior. The blue solid line is the distance between theoretical $G_{1:20}$ and the posterior of sHDP, and the red dotted line is the distance between theoretical $G_{1:20}$ and the posterior of HDP.

Figure 3: The comparison of the successive KL divergences and the accuracy for the phased data set.

distributions, we generate 50 positive integers from each of them representing the component indices which are seen as the observations of the model in order to infer the posterior of $G_{1:20}$. We sample the posterior from both HDP and sHDP. The symmetric KL divergence of successive distributions is shown in Figure 3(a). It can be seen that the symmetric KL values of sHDP are below 1 while that of HDP are much larger, from about 3 to 7. We also compared the distance between the theoretical value of $G_{1:20}$ and the posterior of both models in which the symmetric KL divergence is used as the measure of distance. It can be seen that in the first 10 time intervals, the distances are similar, however, in the last 10 sHDP results much smaller KL distance than that of the HDP.

5.2 Applications on real data set

We parsed the web page of PAMI and collected keywords from approximately 4000 papers published from 1990 to 2015. The keywords from one paper is seen as a document and the years are considered to be the phases. Similarly to [12], we transform each of the documents into a 12 dimensional vector using the method proposed by [10]. We compute the similarity between two documents, which is the number of shared words divided by the total number of words in these two documents. Then we derive 12 eigen vectors for the normalized graph Laplacian [15] and the 12 column matrix of the eigen vectors are seen as the data set. For pre-processing, we change the standard deviation of the columns to be 1.

We assume the likelihood of the data is normally distributed with fixed variance and performed a hierarchical clustering of the data set. The clusters are considered to be the topics of the corpus. In the experiment, we set $\gamma = \alpha = 5$, and the KL bound $B = 3$, and we sample the posterior for 500 iterations. The mixing measures $G_{1:26}$ represents the weights of the keyword clusters for each year.

When we compare the outcomes from both sHDP and HDP, we found that the clusters for each year are similar for both models, but the value of successive symmetric KL divergence can be quite different. The mean of the successive symmetric KL of sHDP is mostly blow 3 while the mean of HDP is around 5. The boxplot of the symmetric KL divergence is shown in Figure 4. Moreover, we track the components weights for the 8 most significant clusters of G_0 through G_1 to G_{26} in figure 5. In this figure, we can see that when the weights of G_0 are large (the first 5 sub-figures in Figure 5), both HDP and sHDP show similar degree of fluctuation in its KL divergence. However, for less significant weight components of G_0, (the last few 3 figures), the smoothing effect of sHDP can be more obviously shown. This illustrates the fact that sHDP can suppress smaller components over times, and hence help achieve smoothness.

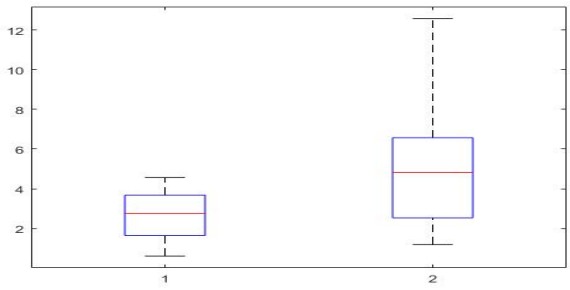

Figure 4: The symmetric KL of successive measures for sHDP and HDP. Left is sHDP and right is HDP.

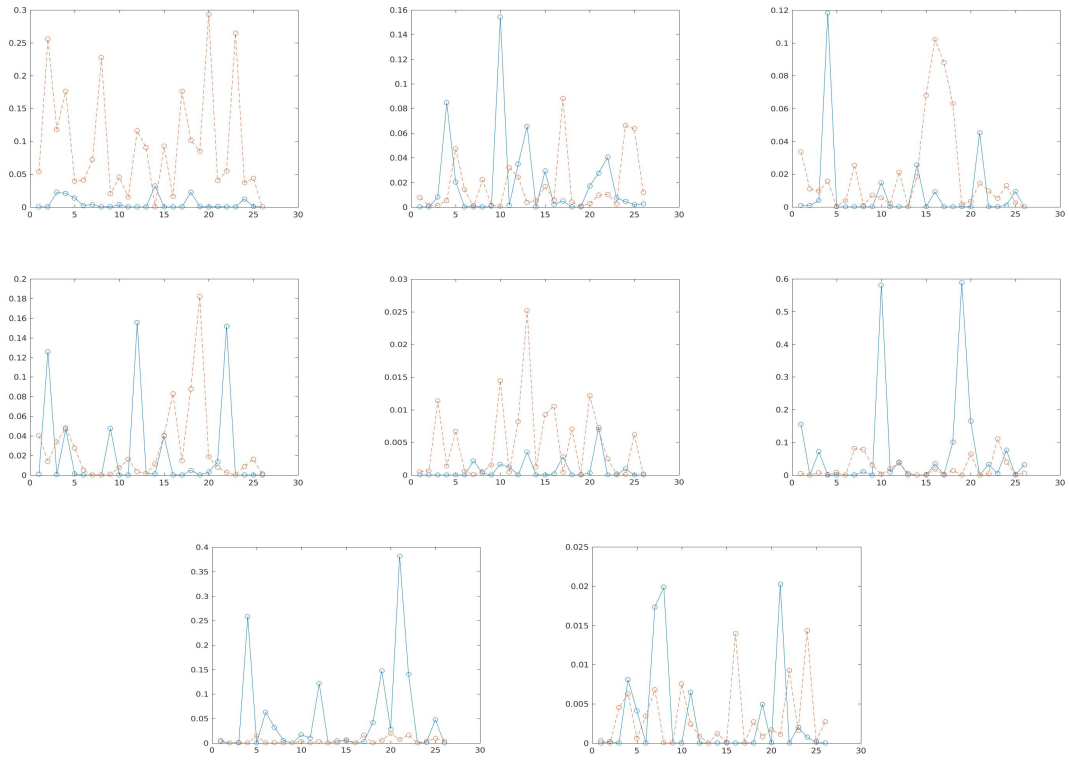

Figure 5: The fluctuation of the weights in G_2. From upper left to lower right are corresponding to the clusters with weights in G_0 in the descending order. The red dotted lines represent the weights of G_2 using sHDP and the blue solid represent the weiths of G_2 using HDP.

6 Conclusion and discussion

We proposed a smoothed Hierarchical Dirichlet Process (sHDP), in which we add a smoothness constraint to the lower layer of a Dirichlet Process (DP), in order to cater to many settings where temporal and/or spatial smoothness between the underlying measures is required. We used symmetric KL divergence for this constraint. Although it is conceptually sound, the traditional symmetric KL in its original form can not be practically applied since the symmetric KL divergence between two countable discrete measures is an infinite sum of real numbers. For this reason, in our paper, we proposed to use aggregated symmetric KL divergence as an alternative measure. We proved that this substitution has the same expectation of the original symmetric KL divergence, and hence can be used appropriately as an alternative bound.

We showed that given a measure at time $t-1$, i.e., G_j, a base measure G_0 and the concentration parameter α, the sampling of the weight of G_{j+1}, i.e., $\beta_{j+1,i}$ can be derived by using a truncated Beta distribution; We show that the truncation boundaries can be solved as roots of a convex function. This has made the solution space easily obtainable.

We show empirically the effect of the bound over the distribution G_{j+1} in figure 2(b): we vary the bound value from 1 to 10 for which we show that the average symmetric KL divergence of the sampled successive distributions have also increased accordingly, which is in line with our expectation. The application of our model to the real dataset shows the smoothness effect of our model, which can be illustrated in both figure 4 and 5. The box plot in 4 show that KL divergence between successive measures, G_j and G_{j+1} of our model is much less than that of the HDP.

We also have analysed the results further to gain some insights into the component-wise smoothness: We choose the eight most significant components of the base measure G_0, and then tracked how its weight changes from G_1 to G_{26} (as we have 26 years of data). The results are shown in figure 5. It is interesting to note that we have a gradual increase in smoothness when its corresponding base measure component is decreased in value.

We also like to point out that our proposed method actually violates de Finetti's theorem about the exchangeable random variables [1]. According to de Finetti's theorem, the prior exists only if the random variables are infinitely exchangeable. However, our assumption on the dependence of successive distributions actually made the prior/base measure G_0 non-existent.

In the current form, the non traceability of the conditional $\mathbb{P}(G_{j+1}|G_j, G_0)$ require us to use particle filter as its inference method. Although it suffices in the settings of our testing dataset which is comprised of only 26 time intervals, it nonetheless can be a computation bottleneck when the application requires us to have much more granular time intervals. Therefore, in our future work, we will experiment with stochastic inference methods for time sequences, similar to that of [13].

References

[1] David J. Ardus. *Exchangeability and related topics*. Springer Berlin Heidelberg, 1985.

[2] Matthew J Beal, Zoubin Ghahramani, and Carl Edward Rasmussen. The Infinite Hidden Markov Model. In *Advances in Neural Information Processing Systems 14*, pages 577–584, 2001.

[3] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3(4-5):993–1022, 2012.

[4] Paul Damien and Stephen G. Walker. Sampling truncated normal, beta, and gamma densities. *Journal of Computational and Graphical Statistics*, 10(2):206–215, 2001.

[5] Arnaud Doucet and Am Johansen. A tutorial on particle filtering and smoothing: fifteen years later. *Handbook of Nonlinear Filtering*, 12:656–704, 2009.

[6] Thomas S. Ferguson. A Bayesian Analysis of Some Nonparametric Problems. *The Annals of Statistics*, 1(2):209–230, 1973.

[7] Emily B Fox. The sticky HDP-HMM: Bayesian nonparametric hidden Markov models with persistent states. *arXiv preprint*, page 60, 2007.

[8] Matt Hoffman, D M Blei, C Wang, and J Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.

[9] Hemant Ishwaran and Lancelot F James. Gibbs Sampling Methods for Stick-Breaking Priors. *Journal of the American Statistical Association*, 96(453):161–173, 2001.

[10] Frbmi Jordan. Learning spectral clustering. *Advances in Neural Information Processing Systems 16 (NIPS)*, pages 305–312, 2004.

[11] Achiim Klenke. *Probability Theory: A Comprehensive Course*. Springer, 2008.

[12] Dahua Lin, Eric Grimson, and John Fisher. Construction of Dependent Dirichlet Processes based on Poisson Processes. In *Advances in Neural Information Processing Systems 23 (Proceedings of NIPS)*, pages 1–9, 2010.

[13] Johnson Matthew and Alan Willsky. Stochastic variational inference for Bayesian time siries models. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1854–1862, 2014.

[14] John Paisley and Lawrence Carin. Hidden markov models with stick-breaking priors. *IEEE Transactions on Signal Processing*, 57(10):3905–3917, 2009.

[15] Max Planck and Ulrike Von Luxburg. A Tutorial on Spectral Clustering A Tutorial on Spectral Clustering. *Statistics and Computing*, 17:395–416, 2006.

[16] Jayaram Sethuraman. A constructive definition of Dirichlet priors. *Statistica Sinica*, 4(2):639–650, 1994.

[17] Yee Whye Teh, Michael I Jordan, Matthew J Beal, and David M Blei. Hierarchical Dirichlet Processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006.

[18] Yee Whye Teh, Kenichi Kurihara, and Max Welling. Collapsed Variational Inference for HDP. In *Advances in neural information processing systems*, 2007.

[19] Chong Wang and David M Blei. Online Variational Inference for the Hierarchical Dirichlet Process. In *International Conference on Artificial Intelligence and Statistics*, volume 2, pages 752–760, 2011.

[20] Mike West. Hyperparameter estimation in Dirichlet process mixture models. *ISDS discussion paper series*, pages #92–03, 1992.